

EventVault

Frequently Asked Questions

© Prism Microsystems, Inc. October 2009

How does EventTracker archive events?

By default, EventTracker does not store event log data in a traditional database; instead Prism stores log data in EventVault, which is Prism's high performance event storage mechanism. EventVault is optimized for the write-once/read many times nature of event log information. In EventVault log data is compressed to less than 10% of the original size, sealed with a SHA-1 checksum and stored in standard CAB files. If 100 million events are archived, a traditional database can grow to 400 GB while EventVault would require approximately 5 GB. When a report is generated, EventTracker automatically selects the required archived data, decompresses and unseals it, and then generates the necessary report. Despite the decompression step, reports via EventVault are still generated faster than using a standard RDBMS and sophisticated caching of the opened event data enabling subsequent report generation to be very fast. The EventVault archives can be stored on any storage device that can be accessed from the EventTracker Manager.

Can you explain more about the archiving process?

EventTracker is designed to receive millions of events daily from hundreds of sources including UNIX events, Windows events, network devices (syslog, pix), databases (log), and flat-files. As events are received from different sources, they are collected in a SQL table as a temporary store. When the temporary table reaches a predetermined optimal size (or on a predefined time interval), the events are compressed in Microsoft cab files, a SHA-1 checksum is calculated and the CAB is moved to a centralized storage location. The size of cab files are tuned to take an advantage of available memory and the read-only behavior of the data store. The ability to store events is limited only by available disk storage.

Why don't you recommend a database for long term archival of events?

A relational database is an excellent choice for managing transactional oriented data – Databases were designed to do that and do it very well. Event data however is not transactional in nature. A relational database is challenged by:

1. **Volume of data:** One million events can represent between 5 and 8GB of data, depending on the size of an event. In a small organization with 12 to 15 domain controllers and with 4 to 5 firewalls, you can easily generate a couple of million of events/day. In a medium size organization it is not uncommon to generate 20 to 30 million events a day, while a large enterprise can generate 50 to 100 million events per day. When you archive even one month of data, this kind of volume

requires an enormous database, a powerful machine to host it, and database admins to maintain it. This will likely not justify your cost/benefit analysis. EventTracker's EventVault makes it much more manageable. Using EventVault the data is compressed (with compression Ratio > 90%) to begin with. Instead of requiring 8GB per 1 million events, EventVault only needs 50MB to store the same 1 million events. There is no database license and no admin time required.

- 2. Performance:** Event data is not a transactional type of data. Events are written once and are never modified. Databases like Oracle and SQL Server consume a huge amount of resources in terms of memory and CPU doing record management in anticipation of the data being modified by multiple users, which never happens with event data. A great deal of their value (and their license fee) is based on this capability. With event data, searching the data within SQL Server or Oracle will be much slower than EventVault because they are not tuned for read-only and unformatted data within an event description. In comparison EventVault is designed for read-only unstructured data. Each EventVault cab file is a small compressed file which can be loaded in memory for a read only search.
- 3. Security:** For compliance and security reasons, organizations require integrity checking against event data and the assurance that the data is not tampered with. With SQL & Oracle this is not possible and realistically it is very hard to protect event data in SQL Server or Oracle against tampering by any competent Administrator. EventVault not only compresses the data, each CAB file is also written with an SHA-1 checksum. This checksum can be verified with the EventVault verification utility which will generate a report detailing if the archives passed or failed the check. In addition you can move the CAB files to an encrypted share, or write once/read only media.
- 4. Cost:** A database license is an expensive investment. When you have a large database you need a large machine that results in a large license cost for the database. When you have a large database, you need an expensive and usually busy database administrator to maintain it for you day to day.

Does the CAB compression affect performance?

No, the CAB files are engineered for performance, scalability and quick read times. In addition there is nothing proprietary about the CAB format. The event storage format is well documented and you can extract event data into any database you want using standard Microsoft tools.

Is EventVault data admissible in court?

Yes. For admissibility in court, you have to prove the report (evidence) is generated from a tamper proof, reliable source. In the process of generating the report, you have to demonstrate the data has not been manipulated by anybody. EventVault is designed such that you can present an EventTracker report as evidence in court. All archived data is written and then sealed with a SHA-1 checksum. Before a report is generated, you can verify and demonstrate that the original data is protected against any tampering and you are generating the report from the original source data.

Is EventVault a proprietary database?

No. It is just Prism's high performance design to achieve unlimited scalability and performance, and there is nothing proprietary about it. Event Data is collected into a SQL table and then compressed into standard Microsoft CAB files. The format of the event storage is well documented and you can extract event data into any database you want using standard Microsoft tools.

How much data can EventVault archive?

EventVault is designed to have unlimited scalability for all practical purposes. It can handle billions of events without significantly compromising the performance of search or report functions. EventVault is made up of thousands of compressed files and each file contains tens of thousands of events. When a report is run, only the files that need to be accessed are accessed.

EventVault is an ideal architecture whether you have hundreds of devices or thousands of devices. There is no limit to the amount of data you can store in the EventVault. In some cases you may want to archive 3 months of events from 10 systems and in others you may want to archive 10 years worth of events from thousands of systems. EventVault handles both.

How do I backup EventVault data?

The backup process is simple. EventVault is made up of a number of compressed files in Microsoft CAB format and are located in a single directory. Each file is appropriately named with create time and dates. All you need to backup is the EventVault directory using your backup software.

Does your reporting engine read the data directly from EventVault?

In many other event log management solutions, report generation from stored data is a laborious affair. You have to load the selected data from offline storage into a SQL database and then tune the database with proper indexing before you can generate the report.

EventTracker does not have concept of offline storage/online database storage. EventVault is always available for report generation any time.

How much system space should I allocate for EventVault?

Though EventVault size is dependent on how many events are generated by each system and is also largely dependent on your audit policy, based on our experience from our 750 plus customers, however, you should plan for an average of about 100MB per server for one year of event archival. If you plan to monitor 100 systems, plan to allocate 10GB for one year.

Can I encrypt EventVault data?

Event data within EventVault is designed with a tamper proof architecture. For most organizations, it will meet their need for security and compliance. You can also use underlying Microsoft technology to encrypt the partition where EventVault data is located.